

Scalable Inference for Markov Processes with Intractable Likelihoods

Jamie Owen*, Darren J. Wilkinson, Colin S. Gillespie

School of Mathematics and Statistics, Newcastle University, UK

Abstract

Bayesian inference for Markov processes has become increasingly relevant in recent years. Problems of this type often have intractable likelihoods and prior knowledge about model rate parameters is often poor. Markov Chain Monte Carlo (MCMC) techniques can lead to exact inference in such models but in practice can suffer performance issues including long burn-in periods and poor mixing. On the other hand approximate Bayesian computation techniques can allow rapid exploration of a large parameter space but yield only approximate posterior distributions. Here we consider the combined use of approximate Bayesian computation (ABC) and MCMC techniques for improved computational efficiency while retaining exact inference on parallel hardware.

Keywords: ABC, particle MCMC, Markov processes, intractable likelihood

1 Introduction

Stochastic kinetic models describe the probabilistic evolution of a dynamical system. Motivated by the need to incorporate intrinsic stochasticity in the underlying mechanics of the systems, stochastic models are increasingly used to model a wide range of real-world problems, particularly, but not limited to, computational systems biology, predator–prey population models, and single species population development (Kitano, 2002; Boys et al, 2008; Gillespie and Golightly, 2010). Systems are governed by a network of reactions which change the state of the system by a discrete amount, and hence are most naturally represented by a continuous time Markov process on a discrete state space. Typically, exact realisations from the model are obtained using the direct method (Gillespie, 1977). There are a number of faster, but approximate algorithms, such as the diffusion approximation, moment closure, or hybrid simulation strategies (Gillespie, 2000; Salis and Kaznessis, 2005; Gillespie, 2009).

Our goal is to perform statistical inference for the parameters that govern the dynamics, where data are partially missing and prior knowledge of reaction rates may be poor. Likelihood functions for such problems are rarely analytically tractable but it is possible to leverage particle Markov chain Monte Carlo (pMCMC) methods to perform exact inference

*j.r.owen@ncl.ac.uk

(Golightly and Wilkinson, 2011; Andrieu et al, 2010). Typically, pMCMC algorithms exhibit a high computational cost, since at each iteration we run a sequential Monte Carlo (SMC) filter, which requires multiple forward simulations. This cost can become overly prohibitive if convergence is slow.

Approximate Bayesian computation (ABC) techniques, allow posterior inference to be made for problems where evaluation of the likelihood function is unavailable, relying heavily on the ability to simulate from the model (Tavare et al, 1997; Pritchard et al, 1999; Beaumont et al, 2002). An approximation to the true posterior distribution of interest is made via samples of parameter vectors, θ , that yield simulated data deemed “close” to the observed data. Here the definition of close is controlled by a tolerance, ϵ , and distance function $\rho(\cdot)$, hence retained samples are from $\pi(\theta|\rho(\mathcal{D}, \mathcal{D}^*) \leq \epsilon)$. The simple rejection algorithm typically leads to high numbers of proposed parameter vectors being rejected for meaningful tolerance values (Pritchard et al, 1999). Further developments within this framework lead to MCMC schemes which avoid calculation of the likelihood, as well as sequential Monte Carlo algorithms, which typically exhibit better acceptance rates than the simple rejection sampler (Marjoram et al, 2003; Del Moral et al, 2006; Sisson et al, 2007; Toni et al, 2009). Such techniques have successfully been applied to stochastic kinetic models for approximate inference (Drovandi and Pettitt, 2011; Fearnhead and Prangle, 2012).

In this article, we explore a scenario where we desire exact posterior inference for model rate parameters, while exploiting parallel hardware. By coupling scalable, approximate and exact techniques, we are able to increase computational efficiency of posterior sampling via parallel particle MCMC and improve convergence properties of chains that suffer expensive burn-in periods to further increase efficiency.

2 Stochastic kinetic models

Consider a network of reactions involving a set of u species $\mathcal{X}_1, \dots, \mathcal{X}_u$ and v reactions $\mathbf{R}_1, \dots, \mathbf{R}_v$ where each reaction \mathbf{R}_i is given by

$$\mathbf{R}_i : \quad p_{i,1}\mathcal{X}_1 + \dots + p_{i,u}\mathcal{X}_u \rightarrow q_{i,1}\mathcal{X}_1 + \dots + q_{i,u}\mathcal{X}_u. \quad (1)$$

Let P be the $v \times u$ matrix of pre-reaction coefficients of reactants $p_{i,j}$, similarly Q , the $v \times u$ matrix of post-reaction coefficients of products. The stoichiometry matrix, defined as

$$S = (Q - P)', \quad (2)$$

is a useful way to encode the structure of the reaction network. Also define X_t to be the vector $(X_{1,t}, \dots, X_{u,t})$ denoting the number of species \mathcal{X} present at time t .

Each reaction \mathbf{R}_i is assumed to have an associated rate constant, θ_i , and a hazard function $h_i(X_t, \theta_i)$ which gives the overall propensity for a type i reaction to occur. The form of $h_i(X_t, \theta_i)$ in many cases can be considered as arising from interactions between components in a well mixed population. This leads to reactions following the law of mass action kinetics where the hazard function for a reaction of type i takes the form

$$h_i = \theta_i \prod_{j=1}^u \binom{X_{j,t}}{p_{i,j}}. \quad (3)$$

If $\theta = (\theta_1, \dots, \theta_v)$ and $h(X_t, \theta) = (h_1(X_t, \theta), \dots, h_v(X_t, \theta))$ then values for θ and X_0 complete full specification of this Markov process.

Algorithm 1 The Direct method

1. Set $t = 0$. Initialise the rate constants θ and initial states X_0 .
2. Calculate the hazard functions $h(X_t, \theta)$ and $h_0(X_t, \theta) = \sum_i^v h_i(X_t, \theta_i)$.
3. Set $t = t + \delta t$ where

$$\delta t \sim \text{Exp}(h_0(X_t, \theta)).$$

4. Simulate the reaction index $j \in (1, \dots, v)$ with probabilities p_j

$$p_j = \frac{h_j(X_t, \theta)}{h_0(X_t, \theta)}.$$

5. Set $X_{t+\delta t} = X_t + S[j]$ where $S[j]$ is the j^{th} column of the stoichiometry matrix S .
 6. If $t < T$ return to 2.
-

2.1 The Direct method

Exact simulations of the evolution of such a system can be obtained via the Direct method (Gillespie, 1977). Algorithm 1 describes the procedure for exact forward simulation of a stochastic kinetic model, given its stoichiometry matrix, S , a vector of reaction rates, θ , an initial state vector X_0 and an associated hazard function $h(X_t, \theta)$. Reactions simulated to occur via the Gillespie algorithm incorporate the discrete and stochastic nature of the system since each reaction, which increments or decrements species levels by discrete amounts, are chosen probabilistically.

Less computationally expensive simulation algorithms such as the chemical Langevin equation (CLE), or Euler Maruyama scheme relax the restriction imposed by the discrete state space but retain the stochasticity of the underlying mechanics, giving approximate realisations of the progression of the species (Gillespie, 2000). For the purpose of exact inference and this article we consider only exact solutions to the stochastic differential equations governing the system. For a more in depth background into stochastic kinetic modelling see Wilkinson (2011).

3 Methods for Bayesian Inference in Intractable Likelihood Problems

Methods for likelihood-free Bayesian inference have become popular due to the increasing use of models in which the likelihood function is either deemed too computationally expensive to compute, or analytically unavailable. In particular, approximate Bayesian computation (ABC), and particle filtering methods have proven to be two classes of algorithms in which one can proceed with inference in a scenario where evaluation of the likelihood is unavailable, commonly termed likelihood-free inference.

Algorithm 2 ABC rejection sampler

1. Generate a candidate parameter vector $\theta^* \sim \pi(\theta)$.
 2. Simulate a candidate data set $\mathcal{D}^* \sim \pi(\mathcal{D}|\theta^*)$.
 3. Calculate a measure of distance between the candidate data, \mathcal{D}^* , and the observed data \mathcal{D} , $\rho(\mathcal{D}, \mathcal{D}^*)$.
 4. Accept θ^* if $\rho(\cdot) < \epsilon$ for some predetermined, fixed, ϵ .
 5. Go to 1.
-

3.1 Approximate Bayesian computation

The goal of approximate Bayesian computation techniques is to obtain a collection of samples from the posterior distribution $\pi(\theta|\mathcal{D})$ where the likelihood function $\pi(\mathcal{D}|\theta)$ is unavailable, due to either high computational cost or analytical intractability. Typically we rely on the assumption that simulation from the model given a set of parameters, θ , is relatively straightforward, as in algorithm 1. Given a set of proposed parameter vectors we would ideally keep any such vectors which yield data simulations that are equal to the observations that we have. In reality however, we will typically match the simulated data to the observations perfectly with low probability. To avoid rejecting all proposals, we instead keep parameters that yield simulations deemed sufficiently close to the observations. We characterise a set of simulated data, \mathcal{D}^* as sufficiently close if, for a given metric, $\rho(\cdot)$, the distance between \mathcal{D}^* and observed data \mathcal{D} is below a tolerance threshold, ϵ . The simple rejection sampler is described in algorithm 2.

Rather than leading to a sample from the exact posterior distribution, the samples instead have the distribution $\pi(\theta|\rho(\mathcal{D}, \mathcal{D}^*) < \epsilon)$. As $\epsilon \rightarrow 0$ this tends to the true posterior distribution if $\rho(\cdot)$ is a properly defined metric on sufficient statistics. Conversely as $\epsilon \rightarrow \infty$ one recovers the prior distribution on the parameters. Since it is usually the case that sufficient statistics are not available in this type of problem, further approximation can be made by choosing a set summary statistics that are not sufficient, but it is hoped that describe the data well. If a large value of ϵ is chosen, a large number of proposed candidates are accepted, but little is learnt about the posterior distribution. Choosing a small tolerance yields a much better approximation to $\pi(\theta|\mathcal{D})$ but at the expense of poorer acceptance rates and hence increased computational cost. This rejection based approach has been included in a number of sampling techniques.

Within the context of this article we consider a sequential approach to ABC inference. An example of such a scheme, based on importance sampling, is described in algorithm 3, see Toni et al (2009) for details.

Algorithm 3 as described here allows us choice of an essentially arbitrary perturbation kernel $K_t(\cdot)$ for preventing degeneracy of particles, and the sequence of tolerances $\epsilon_0, \dots, \epsilon_T$. In practice the choice of $K_t(\cdot)$ will have an important effect on the overall efficiency of the algorithm. A perturbation which gives small moves will typically allow acceptance rates to remain higher, whilst exploring the space slowly. Conversely making large moves allows better exploration of the parameter space, but usually at the cost of reduced acceptance rates. We shall use the ‘optimal’ choice of a random walk perturbation kernel for this algorithm as detailed in Filippi et al (2011). For a multivariate Gaussian random walk

Algorithm 3 Sequential ABC

1. Initialise $\epsilon_0 > \epsilon_1 > \dots > \epsilon_T > 0$ and set the population indicator, $t = 0$.
 2. Set particle indicator, $i = 1$.
 3. If $t = 0$, sample $\theta^{**} \sim \pi(\theta)$
 Else sample θ^* from the previous population $\{\theta_{t-1}^{(i)}\}$ with weights w_{t-1} and perturb to obtain $\theta^{**} \sim K_t(\theta|\theta^*)$
 If $\pi(\theta^{**}) = 0$, return to 3.
 Simulate a candidate dataset $x^* \sim f(x|\theta^{**})$
 If $d(x_0, x^*) \geq \epsilon_t$, return to 3.
 4. Set $\theta_t^{(i)} = \theta^{**}$ and calculate weight for particle $\theta_t^{(i)}, w_t^{(i)}$

$$w_t^{(i)} = \begin{cases} 1, & \text{if } t = 0 \\ \frac{\pi(\theta_t^{(i)})}{\sum_{j=1}^N w_{t-1}^{(j)} K_t(\theta_t^{(j)}, \theta_t^{(i)})}, & \text{if } t > 0 \end{cases}.$$
 If $i < N$ set $i = i + 1$, go to 3
 5. Normalise the weights, if $t < T$, set $t = t + 1$ and go to 2.
-

kernel at population t , we choose $\Sigma^{(t)}$ as

$$\Sigma^{(t)} \approx \sum_{i=1}^N \sum_{k=1}^{N_0} \omega^{(i,t-1)} \tilde{\omega}^{(k)} (\tilde{\theta}^{(k)} - \theta^{(i,t-1)}) (\tilde{\theta}^{(k)} - \theta^{(i,t-1)})' \quad (4)$$

where N is the number of weighted samples we have from the previous population $\pi(\theta|\rho(\mathcal{D}, \mathcal{D}^*) < \epsilon_{t-1})$ with weights ω , N_0 is the number of those samples, $\tilde{\theta}$, that satisfy $\pi(\theta|\rho(\mathcal{D}, \mathcal{D}^*) < \epsilon_t)$ with associated normalised weights $\tilde{\omega}$. This builds on the work of Beaumont et al (2009).

The choice of sequence of tolerances $\epsilon_0, \dots, \epsilon_T$ also has a large effect on the efficiency of the algorithm. A sequence which decreases too slowly will yield high acceptance rates but convergence to the posterior will be slow. Whereas a rapidly decreasing sequence typically yields poor acceptance rates. We will, throughout this article, use an adaptive choice of ϵ_t determined by taking a quantile of the distances at $(t - 1)$.

3.1.1 ABC methods in parallel

Non MCMC-based ABC techniques, such as those described in algorithm 2 and algorithm 3 are often amenable to parallelisation. Since all proposed parameters are independent samples from the same distribution, and the acceptance step for each does not depend on the previously accepted particle, the bulk of computation for a sequential ABC scheme can be run in parallel, greatly reducing overall CPU time required to obtain a sample from the target. Some thread communication is necessary as we move from one distribution in the sequence to the next, something that would be avoided if using the simple rejection sampler, but typically this is small in comparison to the work done in forward simulation.

Coding of a parallel ABC algorithm adds little complexity over a non-parallelised version. This is particularly true for the simple rejection sampler where one effectively

Algorithm 4 Metropolis Hastings

1. Initialise with a random starting value $\theta \sim \pi(\theta)$.
2. Propose a move to a new candidate $\theta^* \sim q(\theta^*|\theta)$.
3. Accept the move with probability

$$\min \left\{ 1, \frac{\pi(\mathcal{D}|\theta^*)\pi(\theta^*)q(\theta|\theta^*)}{\pi(\mathcal{D}|\theta)\pi(\theta)q(\theta^*|\theta)} \right\}, \quad (5)$$

else remain at θ .

4. Return to 2.
-

just runs a sampler on each of N processors and no communication between threads is needed.

3.2 Markov Chain Monte Carlo

Suppose interest lies in $\pi(\theta|\mathcal{D})$, and that we wish to construct a MCMC algorithm whose stationary distribution is exactly this posterior. Using an appropriate proposal function $q(\theta^*|\theta)$ we can construct a Metropolis Hastings algorithm to do this as described in algorithm 4. This can often be impractical due to unavailability of the likelihood term, $\pi(\mathcal{D}|\theta)$.

In practice one normally chooses a proposal function $q(\theta^*|\theta)$ such that proposed moves from θ to θ^* are distributed symmetrically about θ giving rise to random walk Metropolis Hastings. Typical choices include a uniform $\mathcal{U}(\theta - \sigma, \theta + \sigma)$ or Gaussian $\mathcal{N}(\theta, \Sigma)$ distribution. In this article we shall consider Gaussian innovations for random walk proposals. It has been shown that under various assumptions about the target, the optimal scaling for a random walk Metropolis algorithm using a Gaussian proposal, Σ_q , for a d dimensional distribution is

$$\Sigma_q = \frac{(2.38)^2}{d} \Sigma, \quad (6)$$

where Σ is the posterior covariance (Roberts et al, 1997; Roberts and Rosenthal, 2001). It is of note however that typically the posterior covariance is unavailable.

3.2.1 Parallel MCMC

MCMC algorithms are somewhat less amenable to parallelisation than rejection sampling algorithms. Essentially there are two options to be considered, parallelisation of a single MCMC chain or the construction of parallel chains, for an in depth discussion see Wilkinson (2006). Parallelisation of a single chain in many scenarios is somewhat difficult due to the inherently iterative nature of a Markov chain sampler and will not be discussed in detail here. Running parallel chains, however, is straightforward and potentially useful.

3.2.2 Parallel chains

In practice, chains initialised with an arbitrary starting point will be, after an appropriate burn-in period, essentially sampling from the target. In the context of a serial implementation there is still some debate over whether it is better to run a single long chain, or many shorter chains. The benefits of the single chain are that any burn in period must only be suffered once, where as the argument for multiple shorter chains is that one may better diagnose convergence to the stationary distribution. The argument changes when considering parallel implementation. Indeed, if burn-in periods are relatively short, running independent chains on separate processors can be a very time efficient way of learning about the distribution of interest.

Burn-in is still a potential limiter on the scaling of the performance to be had when employing parallel chains with the number of processors. The greater the burn-in period of a chain, the more time each processor has to waste computing samples that will eventually be thrown away. The theoretical speed up calculation given N processors to obtain n stored samples with a burn-in period b is

$$\text{Speed-up}(N) = \frac{b + n}{b + \frac{n}{N}}, \quad (7)$$

which is clearly limited for any $b > 0$, as $N \rightarrow \infty$ (Wilkinson, 2006). A “perfect” parallelisation of multiple chains then is one in which we have no burn-in period, i.e initialise each chain with an independent draw from the target. The performance gain when run on N processors is then of factor N . Clearly this “perfect” situation will not be possible to implement in practice since we typically are unable to sample from the target.

3.2.3 Particle MCMC

Re-consider the Metropolis Hastings MCMC algorithm from section 3.2. The ability to construct such a scheme relies on the evaluation of the likelihood function $\pi(\mathcal{D}|\theta)$ among others. Crucially, the problems of interest discussed in this article are such that the likelihood term $\pi(\mathcal{D}|\theta)$ is unavailable. However Andrieu and Roberts (2009) proposed a pseudo marginal MCMC approach to this problem. In the acceptance ratio of the MCMC scheme, replace the intractable likelihood, $\pi(\mathcal{D}|\theta)$, with a Monte Carlo estimate. It can be shown that, provided $E[\hat{\pi}(\mathcal{D}|\theta)] = \pi(\mathcal{D}|\theta)$, the stationary distribution of the resulting Markov chain is exactly the distribution of interest. In the context of inference for state space models, it is natural to make use of sequential Monte Carlo techniques through a bootstrap particle filter, (Doucet et al, 2001), to obtain estimates of the likelihood $\hat{\pi}(\mathcal{D}|\theta)$. The bootstrap particle filter for application to Markov processes is described in algorithm 5.

Substituting the estimate of $\hat{\pi}(\mathcal{D}|\theta)$ in place of the likelihood function yields a MH algorithm with acceptance probability

$$\min \left\{ 1, \frac{\hat{\pi}(\mathcal{D}|\theta^*)\pi(\theta^*)q(\theta|\theta^*)}{\hat{\pi}(\mathcal{D}|\theta)\pi(\theta)q(\theta^*|\theta)} \right\}. \quad (8)$$

It turns out that the particle filter’s estimate of marginal likelihood is unbiased, giving rise to an “exact approximate” pseudo-marginal MCMC algorithm. This use of a SMC sampler within a pseudo-marginal algorithm is a special case of the particle marginal Metropolis Hastings (PMMH) algorithm described in Andrieu et al (2010), which in general can be

Algorithm 5 The boot-strap particle filter

At time t we have a set of N particles $\mathbf{X}_t^* = \{(x_t^i, \pi_t^i) : i = 1, \dots, N\}$. The filter assumes fixed parameters and so we drop the θ from our notation. $t \in (0, 1, \dots, T)$ for observed data with T discrete time-point observations.

1. Initialise at \mathbf{X}_0^* , a set of N independent draws from our prior distribution on the state.
2. At time t , suppose we have a sample $\mathbf{X}_t^* \sim \pi(\mathbf{X}_t | D_{1:t})$.
3. Sample a set of indices for candidates for forward simulation, I_t^i according to the weights π_t .
4. Simulate forward from the model the chosen paths, $x_{t+1}^i \sim \pi(x_{t+1}^i | x_t^{I_t^i})$.
5. Calculate weights, $w_{t+1}^i = p(d_{t+1} | x_{t+1}^i)$, and normalise to set $\pi_{t+1}^i = \frac{w_{t+1}^i}{\sum_{j=1}^N w_{t+1}^j}$.

Define $\hat{p}(d_t | D_{1:t-1}) = \frac{1}{N} \sum_{i=1}^N w_t^i$, then $\hat{p}(D_{1:t}) = \prod_{t=1}^T \hat{p}(d_t, D_{1:t-1})$.

used to target the full joint posterior on the state and rate $\pi(\theta, X | \mathcal{D})$. In this article we consider only the pseudo-marginal special case. As discussed in section 3.2.1 this type of scheme does not lend itself well to parallelisation unless we have the opportunity to initialise chains with a sample from the posterior distribution $\pi(\theta | \mathcal{D})$. We recognise that rather sophisticated parallel particle filter algorithms which have the potential for increasing the speed of moves within a single chain are available. However the speed up is limited by necessary communication and synchronisation between processors and typically do not scale well.

The pseudo-marginal MCMC scheme requires specification of a number of particles, N , to be used in the boot-strap filter. It was noted, (Andrieu and Roberts, 2009), that the efficiency of the scheme decreases as the variance of the estimated marginal likelihood increases. This can be overcome by increasing the value of N , albeit at an increased computational cost. An optimal choice of N was the subject of Pitt et al (2012) and Doucet et al (2012). The former suggest that N should be chosen such that the variance in the estimated log-likelihood is around 1 while the latter show that the penalty is small for values between 0.25 and 2.25.

3.3 Combined use of ABC and pMCMC

Since it is typically not possible to initialise a MCMC chain with a draw from the desired target, we propose an approach to parallel MCMC by choosing initial parameter vectors according to samples from an approximate posterior distribution. The intuition is that if we have a reasonable approximation to the target of interest, samples from the approximation will closely match those from $\pi(\theta | \mathcal{D})$. Because of this we expect that after a very short burn-in we are sampling from the desired target. Hence we are approaching the scenario of perfect parallelisation of MCMC. Clearly the better the approximation the shorter the burn-in for each chain will be.

If we first run an ABC scheme targeting an approximation to this distribution, we can exploit parallel hardware to eliminate a large region of prior parameter space very quickly.

Then take a set of N independent samples from $\pi(\theta|\rho(\mathcal{D}, \mathcal{D}^*) < \epsilon)$ as initialisations to N independent, parallel, pMCMC chains each targeting the exact posterior distribution $\pi(\theta|\mathcal{D})$.

We can in some sense consider this process of obtaining a sample from an ABC posterior as an artificial burn-in period. Crucially however, since the ABC algorithm yields a large set of samples from $\pi(\theta|\rho(\mathcal{D}, \mathcal{D}^*) \leq \epsilon)$ the computational price of performing the artificial burn-in has only to be paid once and can be parallelised. With a set of samples we can initialise an arbitrary number of MCMC chains each with an independent parameter vector θ_0 which comes from a distribution close to the target $\pi(\theta|\mathcal{D})$.

Since ABC is being used only for initialisation, it is not used as a prior for MCMC algorithm, nor does it form part of the proposal. Hence, the ABC approximation does not affect the exactness of the pMCMC target.

3.4 Random Walk pMCMC using ABC

We now refer back to the optimal choice of Gaussian random walk kernel of Roberts and Rosenthal (2001) mentioned in section 3.2. Since we hope that we start with a reasonable approximation to the true posterior, we likewise consider that the covariance of a sample from $\pi(\theta|\rho(\mathcal{D}, \mathcal{D}^*) < \epsilon)$ denoted here Σ_{ABC} , will be close to the covariance of true posterior samples. This allows us to use the approximate posterior covariance structure as an informative tool for calibrating the random walk innovations. In most MCMC applications, since the true posterior variance is unknown, Σ_q requires tuning, typically from pilot runs. In the case of using Σ_{ABC} we hope that we can also remove the necessity to tune Σ_q . In practice we take

$$\Sigma_q = \frac{(2.38)^2}{d} \Sigma_{\text{ABC}}. \quad (9)$$

In addition we use our approximation to automatically tune the number of particles required for the particle filter. In practice we take the average of the approximate distribution, $\bar{\theta}_{\text{ABC}}$ to calculate the number of particles required to satisfy

$$\text{Var}(\hat{\pi}(\mathcal{D}|\bar{\theta}_{\text{ABC}})) \simeq 2, \quad (10)$$

in line with Sherlock et al (2013). The hybrid ABC pMCMC algorithm is outlined in algorithm 6.

4 Applications

4.1 The Lotka–Volterra system

The Lotka–Volterra predator–prey system is a simple stochastic kinetic model (Lotka, 1925; Volterra, 1926). The system is characterised by a set of three reactions detailing interactions between two species; prey \mathcal{X}_1 and predators \mathcal{X}_2 :



The reactions \mathbf{R}_1 , \mathbf{R}_2 and \mathbf{R}_3 can be thought of as prey birth, an interaction resulting in a prey death and predator birth, and predator death respectively. Whilst this is a

Algorithm 6 Hybrid ABC pMCMC

1. Run an ABC algorithm targeting $\pi(\theta|\rho(\mathcal{D}, \mathcal{D}^*) \leq \epsilon_T)$
 2. Initialise multiple MCMC chain with a sample $\theta_0 \sim \pi(\theta|\rho(\mathcal{D}, \mathcal{D}^*) \leq \epsilon_T)$ and set $i = 1$
 3. Propose $\theta^* \sim q(\theta^*|\theta)$ for $q(\cdot)$ a Gaussian random walk kernel, $\mu_q = \theta$, $\Sigma_q = \frac{(2.38)^2}{d} \Sigma_{\text{ABC}}$
 4. Approximate $\pi(\mathcal{D}|\theta^*)$ via a bootstrap particle filter, $\hat{\pi}(\mathcal{D}|\theta^*)$ where the number of particles is chosen as in (10)
 5. With probability
$$\min \left\{ 1, \frac{\hat{\pi}(\mathcal{D}|\theta^*)\pi(\theta^*)q(\theta|\theta^*)}{\hat{\pi}(\mathcal{D}|\theta)\pi(\theta)q(\theta^*|\theta)} \right\}$$
set $\theta^{(i)} = \theta^*$ else set $\theta^{(i)} = \theta^{(i-1)}$
 6. Set $i := i + 1$ and return to 3.
-

relatively simple example of a model in this context, it highlights many of the difficulties associated with more complex systems. We summarise the system by its stoichiometry matrix, S and hazard function $h(\mathbf{X}, \theta)$:

$$S = \begin{pmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \end{pmatrix}, \quad h(\mathbf{X}, \theta) = (\theta_1 X_1, \theta_2 X_1 X_2, \theta_3 X_2). \quad (12)$$

Figure 1(a) is a set of synthetic, noisy observations of the two species simulated using the Gillespie algorithm with true parameter values $\log(\theta) = (0, -5.30, -0.51)$ together with a Gaussian measurement error structure, $d_{j,t} \sim \mathcal{N}(X_{j,t}, 10^2)$.

4.2 pMCMC for Lotka–Volterra

Exact parameter inference is possible for this system using pseudo marginal MCMC pMCMC scheme; provided the chain is initialised near the posterior mode (Wilkinson, 2011). However under poor initialisation, a pseudo-marginal particle MCMC scheme will perform very badly.

Suppose we have little prior knowledge on the rate parameters,

$$\log(\theta_i) \sim \mathcal{U}(-8, 8), \quad i = 1, 2, 3. \quad (13)$$

We take a prior distribution for the initial state X_0 on each individual species as Poisson distributions with rate parameter equal to the true initial conditions,

$$x_{1,0} \sim \text{Pois}(50), \quad x_{2,0} \sim \text{Pois}(100). \quad (14)$$

Further we assume that the variance, σ^2 , of the Gaussian measurement error is known.

Using a Gaussian random walk proposal $q(\theta^*|\theta)$ we construct a MH algorithm using a bootstrap particle filter, targeting $\pi(\theta|\mathcal{D})$ as in algorithm 4. Figure 1(b) shows that when initialising the chain using random draws from the weakly informative prior, the chains

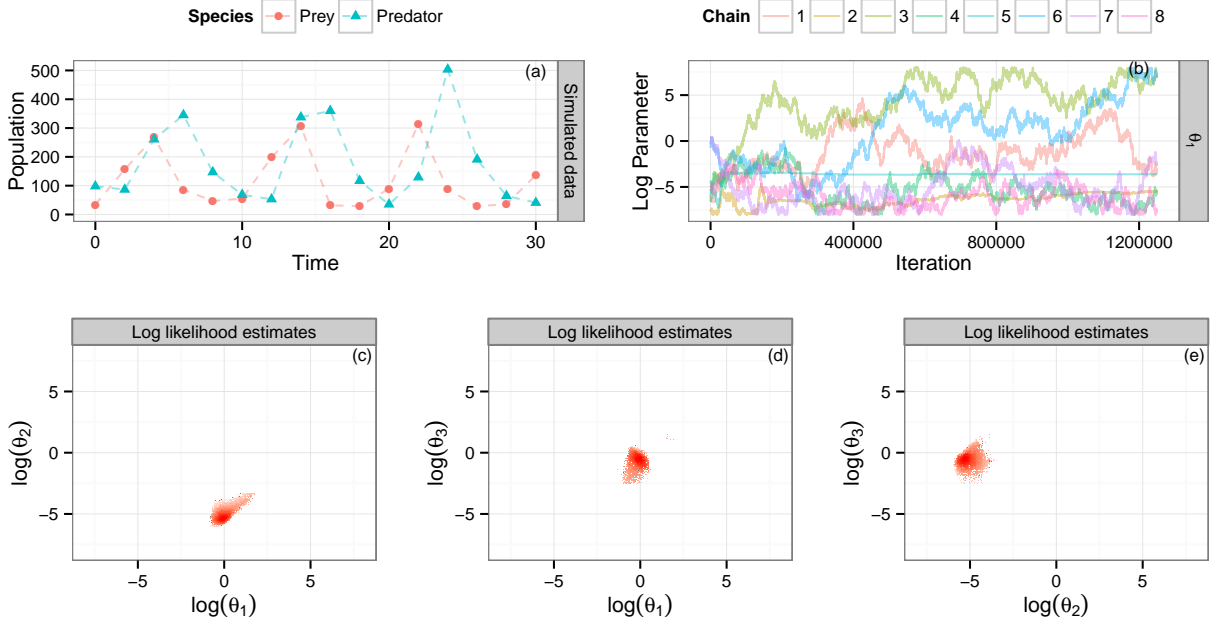


Figure 1: Investigating computational issues with pMCMC for the Lotka–Volterra model. (a) The true underlying synthetic data set. Species are observed at discrete time points and corrupted with $N(0, 10^2)$ noise. (b) Eight trace plots of θ_1 from pMCMC chains initialised with random draws from the prior (see 13). Due to numerical issues, the chains wander the space randomly. (c,d,e) are heat-maps of log-likelihood estimates from the bootstrap particle filter; the white areas is where the estimated log-likelihood is negative infinity. In each of these plots the missing parameter is taken to be fixed at the true value

wander the space at random failing to converge due to the numerical underflow issue. Further investigation shows why this is happening. Figure 1 (c,d,e) show log-likelihood estimates $\hat{\pi}(\mathcal{D}|\theta)$ using a bootstrap particle filter as a heat map. The plot shows that the likelihood surface has a very steep gradient around the true values, and that a large proportion of the space leads to an estimate of negative infinity. This has two consequences, the first being that when we are a reasonable distance from a region of high likelihood, to available numerical accuracy, we obtain estimates that are numerically equal to negative infinity. This means that exploration of the space depends only on the random walk moves, unguided by the likelihood structure. On top of this, small proposed moves can lead to large variation in the estimated likelihood, leading to poor exploration of the tails in the posterior distribution. Whilst we are guaranteed to eventually converge to the stationary distribution, the required computational cost, without carefully thought out initialisation, could be very high. We note that this is not a failure of the theory or algorithm, but a consequence of the computational accuracy available.

We therefore apply the proposed ABC initialisation for a “perfect” parallel pMCMC scheme.

4.2.1 Results for Lotka–Volterra model

We use a sequence of seven distributions $\pi(\theta|\rho(\mathcal{D}, \mathcal{D}^*) < \epsilon_t)$, $t = 0, \dots, 6$ to obtain our approximation to the posterior. For each population t we take ϵ_t as the 0.3 quantile of the distribution of distances from the samples at $t - 1$ and propose candidates until we reach a sample of 1000 for each t . At each stage we perform the forward simulation of the model

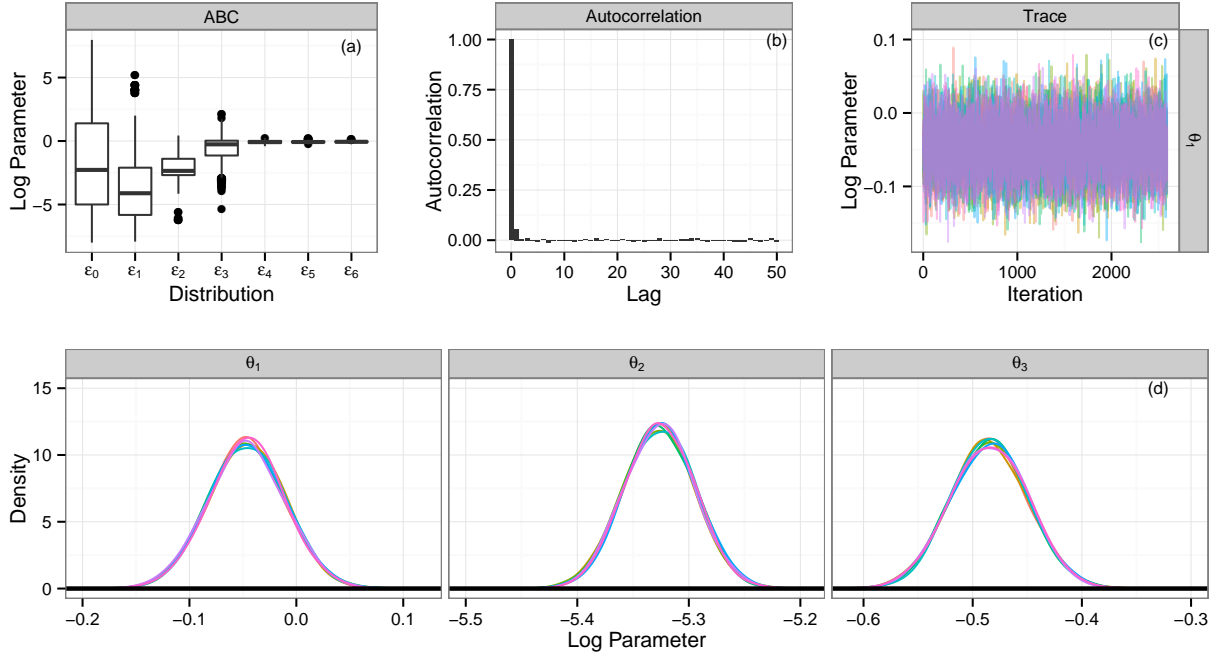


Figure 2: Analysis of results for the synthetic data for the Lotka–Volterra model. (a) are successive distributions of $\log(\theta_1)$ in the sequential ABC scheme, algorithm 3. (b) and (c) are analysis of eight parallel MCMC chains for $\log(\theta_1)$. Note that each chain is sampling from the same stationary distribution and mixing appears good. (d) are the posterior densities for $\log(\theta)$, each chain leads to a posterior density plot that is very close to that of every other chain. True values $\log(\theta) = (0.0, -5.30, -0.51)$ are well identified

given proposed parameter values in parallel. Figure 2 (a) shows summaries for the reaction 1 rate parameter θ_1 for each of the distributions in the series. The distributions quickly remove potential initialisation points and converge around the true value $\log(\theta_1) = 0$. Given the sample from $\pi(\theta|\rho(\mathcal{D}, \mathcal{D}^*) < \epsilon_7)$ we initialise eight MCMC chains on separate processors with random draws. The results in figure 2(b,c,d) are then, the 20,000 pooled samples from the eight independent parallel chains, each of which has been thinned by a factor of 100 to give 2,500 samples. Each chain is sampling from the same stationary distribution as seen in the trace plot for θ_1 , figure 2 (c), and mixing is good, figure 2 (b). Further the true parameter values, $\log(\theta) = (0, -5.30, -0.51)$, used to simulate the data are well identified within the posterior densities, figure 2 (d).

4.3 Real-data problem – aphid model

Next we consider a model of aphid population dynamics as proposed in Matis et al (2007). The system can be represented by the following reactions:



and summarised in the same way as with the Lotka–Volterra example in terms of its stoichiometry matrix S , and hazard function $h(\mathbf{X}_t, \theta)$,

$$S = \begin{pmatrix} 1 & -1 \\ 1 & 0 \end{pmatrix}, \quad h(\mathbf{X}, \theta) = (\lambda N, \mu NC), \tag{16}$$

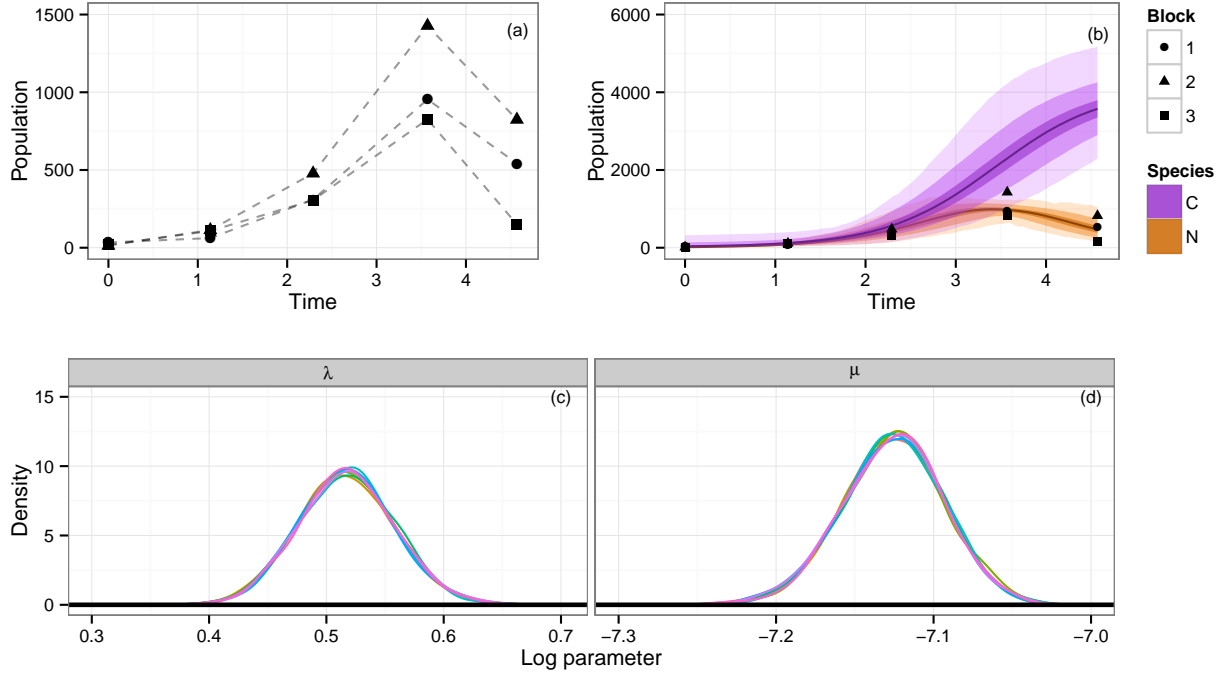


Figure 3: Analysis of the real data for the aphid growth model. (a) There are the three data sets of aphid counts each with of five observations. (b) The posterior predictive model fit given a sample from the collection of posterior densities. (c,d) Output from each MCMC chain, highlight that we are sampling from the same stationary distribution

$$\mathbf{X} = (N, C), \theta = (\lambda, \mu).$$

N_t and C_t are the numbers of aphids alive at time t and the cumulative number of aphids that have lived up until time t respectively. In the first reaction, when a new aphid is introduced into the system we get an increase in both the current number of aphids and the cumulative count. When an aphid is removed the number of aphids decreases by one but the cumulative count remains the same. Aphid death is considered to be proportional not only to the number of aphids currently in the system, but also to the cumulative count, representing the idea that over time they are making their own environment less inhabitable.

Given some initial conditions $\mathbf{X}_0 = (N_0, C_0)$ and a set rate of parameters θ , we can simulate from the model much like with any stochastic system of this type. However the difficulty in practice with the model presented here is that we will never have observations of the cumulative count. Observations then are noisy, discrete time measurements of just a single species of the system.

4.3.1 Aphid data

We now consider the data described in Matis et al (2008) consisting of cotton-aphid counts for twenty seven treatment-block combinations. The treatments consisted of three nitrogen levels(blanket, variable and none), three irrigation levels(low, medium and high) and three blocks. The sampling times of the data are $t = 0, 1.14, 2.29, 3.57, 4.57$ weeks, or every seven to eight days. We restrict our investigation to a single treatment combination, three data sets with blanket nitrogen level and low irrigation. If we denote the block by $i \in \{1, 2, 3\}$ then the data \mathcal{D}_i is the number of aphids, N , in block i at each time t . The data are plotted in figure 3(a).

We make the assumption that the counts are observed with error such that

$$d_t \sim \text{Pois}(x_t). \quad (17)$$

We again use a set of weakly informative priors on the rate parameters θ

$$\log(\theta_i) \sim \mathcal{U}(-8, 8), \quad i = 1, 2. \quad (18)$$

We place a prior of the form

$$C_0 = N_0 + g, \quad g \sim \text{Geom}(0.03). \quad (19)$$

We treat the three sets of observations as repeats of the same experiment. A full treatment of all twenty-seven data sets using a fixed effects model can be found in Gillespie and Golightly (2010). We consider the initial aphid counts to be the true values, as in Gillespie and Golightly (2010), on the basis that there should be no error in counting such small populations.

4.4 Results for the aphid growth model

We use the same criteria for the choice of ϵ_t for the ABC section of the inference as with the Lotka–Volterra model in 4.2.1, namely the 0.3 quantile of the distribution of distances. A sequence of five distributions gives us 1000 samples from $\pi(\theta|\rho(\mathcal{D}, \mathcal{D}^*) < \epsilon_5)$ with which we, as before, initialise eight parallel chains. We record 20,000 samples from the exact target posterior $\pi(\theta|\mathcal{D})$ after appropriate thinning. Figure 3(c and d) shows the analysis of the MCMC chains. Again we find that each chain is sampling from the same target and posterior densities are very close from all eight chains. Figure 3(b) shows posterior predictive quantiles given a sample from the pooled posterior distribution. Model fit appears to be reasonable.

5 Discussion

We have proposed an approach to exact inference for Markov processes that combines the relative strengths of ABC and pMCMC methodology to increase computational efficiency through use of parallel hardware. Through use of an approximation to the posterior distribution of interest, obtained via a sequential ABC algorithm which is easy to parallelise, we can set up a parallel implementation of pMCMC which has numerous desirable properties. By enabling the construction of independent parallel chains initialised close to the stationary distribution, this enables fast convergence and sampling from an exact posterior distribution that scales well with available computational resources. Algorithmic tuning parameters required for pMCMC, such as the variance of Gaussian random walk proposals and numbers of particles for the particle filter can be chosen without the need for additional pilot runs, as a consequence of having a sample from an ABC posterior. In addition, independent parallel chains allow verification of convergence and the computational saving in burn-in times extends to repeat MCMC analyses.

We have demonstrated this approach by applying it to two stochastic kinetic models. With the Lotka–Volterra predator prey system, a relatively simple model in which both species can be observed, we highlighted clear issues with practical implementation of a pseudo-marginal approach in a scenario in which prior information on reaction rate

parameters is poor. This issue can be alleviated by first obtaining a sample from an approximation to the posterior, then using it to guide an exact pMCMC scheme. The approach discussed performed similarly well in the application to a set of real data for a model for aphid growth dynamics in which one of the species in the system can never be observed, where we again imposed weak prior conditions on the rate parameters governing the system and had access to repeat data.

References

- Andrieu C, Roberts GO (2009) The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics* 37(2):697–725
- Andrieu C, Doucet A, Holenstein R (2010) Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72(3):269–342
- Beaumont MA, Zhang W, Balding DJ (2002) Approximate Bayesian computation in population genetics. *Genetics* 162(4):2025–2035
- Beaumont MA, Cornuet JM, Marin JM, Robert CP (2009) Adaptive approximate Bayesian computation. *Biometrika* 96(4):983–990
- Boys RJ, Wilkinson DJ, Kirkwood TB (2008) Bayesian inference for a discretely observed stochastic kinetic model. *Statistics and Computing* 18(2):125–135
- Del Moral P, Doucet A, Jasra A (2006) Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68(3):411–436
- Doucet A, de Freitas N, Gordon N (2001) Sequential Monte Carlo methods in practice
- Doucet A, Pitt M, Kohn R (2012) Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *arXiv:12101871*
- Drovandi CC, Pettitt AN (2011) Estimation of parameters for macroparasite population evolution using approximate Bayesian computation. *Biometrics* 67(1):225–233
- Fearnhead P, Prangle D (2012) Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 74(3):419–474
- Filippi S, Barnes CP, Cornebise J, Stumpf MP (2011) On optimality of kernels for approximate Bayesian computation using sequential Monte Carlo. *Statistical Applications in Genetics and Molecular Biology* pp 1–21
- Gillespie CS (2009) Moment-closure approximations for mass-action models. *IET Systems Biology* 3(1):52–8
- Gillespie CS, Golightly A (2010) Bayesian inference for generalized stochastic population growth models with application to aphids. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 59(2):341–357

- Gillespie DT (1977) Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry* 81(25):2340–2361
- Gillespie DT (2000) The chemical Langevin equation. *The Journal of Chemical Physics* 113(1):297–306
- Golightly A, Wilkinson DJ (2011) Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus* 1(6):807–820
- Kitano H (2002) Computational systems biology. *Nature* 420(6912):206–210
- Lotka AJ (1925) *Elements of physical biology*. Williams & Wilkins Baltimore
- Marjoram P, Molitor J, Plagnol V, Tavaré S (2003) Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences* 100(26):15,324–15,328
- Matis JH, Kiffe TR, Matis TI, Stevenson DE (2007) Stochastic modeling of aphid population growth with nonlinear, power-law dynamics. *Mathematical Biosciences* 208(2):469–494
- Matis TI, Parajulee MN, Matis JH, Shrestha RB (2008) A mechanistic model based analysis of cotton aphid population dynamics data. *Agricultural and Forest Entomology* 10(4):355–362
- Pitt MK, Silva RdS, Giordani P, Kohn R (2012) On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *Journal of Econometrics* 171(2):134–151
- Pritchard JK, Seielstad MT, Perez-Lezaun A, Feldman MW (1999) Population growth of human y chromosomes: a study of y chromosome microsatellites. *Molecular Biology and Evolution* 16(12):1791–1798
- Roberts GO, Rosenthal JS (2001) Optimal scaling for various Metropolis-Hastings algorithms. *Statistical science* 16(4):351–367
- Roberts GO, Gelman A, Gilks WR (1997) Weak convergence and optimal scaling of random walk metropolis algorithms. *The Annals of Applied Probability* 7(1):110–120
- Salis H, Kaznessis Y (2005) Accurate hybrid stochastic simulation of a system of coupled chemical or biochemical reactions. *Journal of Chemical Physics* 122:054,103
- Sherlock C, Thiery AH, Roberts GO, Rosenthal JS (2013) On the efficiency of pseudo-marginal random walk Metropolis algorithms. *arXiv:13097209*
- Sisson S, Fan Y, Tanaka MM (2007) Sequential Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences* 104(6):1760–1765
- Tavare S, Balding DJ, Griffiths R, Donnelly P (1997) Inferring coalescence times from dna sequence data. *Genetics* 145(2):505–518
- Toni T, Welch D, Strelkowa N, Ipsen A, Stumpf MP (2009) Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface* 6(31):187–202

- Volterra V (1926) Fluctuations in the abundance of a species considered mathematically. *Nature* 118:558–560
- Wilkinson DJ (2006) Parallel Bayesian computation. *Statistics Textbooks and Monographs* 184:477
- Wilkinson DJ (2011) *Stochastic modelling for systems biology*, vol 44. CRC press